

AD-A204 680

FILE COPY

4

**VARIETIES OF LEARNING
IN SOAR: 1987**

Technical Report AIP-6

D.M. Steier, J.E. Laird, A. Newell,
P.S. Rosenbloom, R.A. Flynn, A. Golding,
T.A. Polk, O.G. Shivers
A. Unruh, G.R. Yost

Carnegie-Mellon University
University of Michigan
Stanford University

**The Artificial Intelligence
and Psychology Project**

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
S DEC 27 1988 **D**
E

Approved for public release; distribution unlimited.

88 12 27 178

4

VARIETIES OF LEARNING IN SOAR: 1987

Technical Report AIP-6

D.M. Steier, J.E. Laird, A. Newell,
P.S. Rosenbloom, R.A. Flynn, A. Golding,
T.A. Polk, O.G. Shivers
A. Unruh, G.R. Yost

Carnegie-Mellon University
University of Michigan
Stanford University

29 September 1987

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



DTIC
ELECTE
DEC 27 1988
S D
E

This research was supported by the Computer Sciences Division, Office of Naval Research and DARPA under Contract Number N00014-86-K-0678; and by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under contracts F33615-81-K-1539 and N00039-83-C-0136. Additional partial support was provided by the Sloan Foundation and some computing support was supplied by the SUMEX-AIM facility (NIH grant number RR-00785). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Sloan Foundation, the National Institute of Health, or the U. S. Government. Reproduction in whole or in part is permitted for purposes of the the United States Government. Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 6			5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University		6b. OFFICE SYMBOL (If applicable)		7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research (Code 1113)
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (If applicable)		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS p400005up2017-4-86	
			PROGRAM ELEMENT NO N/A	PROJECT NO N/A
			TASK NO N/A	WORK UNIT ACCESSION NO N/A
11. TITLE (Include Security Classification) Varieties of Learning in Soar: 1987				
12. PERSONAL AUTHOR(S) D.M. Steir, J.E. Laird, A. Newell, P.S. Rosenbloom, R.A. Flynn, A. Golding, T.A. Polk, O.G. Shivers, A. Unruh, G.R. Yost				
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 86Sept15 to 91Sept14		14. DATE OF REPORT (Year, Month, Day) 87 September 29
15. PAGE COUNT 13				
16. SUPPLEMENTARY NOTATION				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	Artificial Intelligence, Machine learning, cognitive architecture, (SDW)	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)				
<p>→ Soar is an architecture for intelligence that integrates learning into all of its problem-solving behavior. The learning mechanism, chunking, has been studied experimentally in a broad range of tasks and situations. This paper summarizes the research on chunking in Soar, covering the effects of chunking in different tasks, task-independent applications of chunking and our theoretical analyses of effects and limits of chunking. We discuss what and when Soar has been able to learn so far. The results demonstrate that the variety of learning in Soar arises from variety in problem solving, rather than from variety in architectural mechanisms. <i>Keywords:</i></p>				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz			22b. TELEPHONE (Include Area Code) (202) 696-4302	22c. OFFICE SYMBOL N00014

Varieties of Learning in Soar: 1987

David M. Steier¹

John E. Laird²

Allen Newell¹

Paul S. Rosenbloom³

Rex A. Flynn¹

Andrew Golding³

Thad A. Polk¹

Olin G. Shivers¹

Amy Unruh³

Gregg R. Yost¹

(STEIER@C.CS.CMU.EDU)

(LAIRD@CAEN.ENGIN.UMICH.EDU)

(NEWELL@C.CS.CMU.EDU)

(ROSENBLROOM@SUMEX-AIM.STANFORD.EDU)

(RAF@F.GP.CS.CMU.EDU)

(GOLDING@SUMEX-AIM.STANFORD.EDU)

(TAP@F.GP.CS.CMU.EDU)

(SHIVERS@H.CS.CMU.EDU)

(UNRUH@SUMEX-AIM.STANFORD.EDU)

(GRY@G.CS.CMU.EDU)

¹Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA 15208 U.S.A

²Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 U.S.A.

³Knowledge Systems Laboratory, Department of Computer Science, Stanford University, 701 Welch Road (Bldg. C), Palo Alto, CA 94304 U.S.A.

Abstract

Soar is an architecture for intelligence that integrates learning into all of its problem-solving behavior. The learning mechanism, chunking, has been studied experimentally in a broad range of tasks and situations. This paper summarizes the research on chunking in Soar, covering the effects of chunking in different tasks, task-independent applications of chunking and our theoretical analyses of effects and limits of chunking. We discuss what and when Soar has been able to learn so far. The results demonstrate that the variety of learning in Soar arises from variety in problem solving, rather than from variety in architectural mechanisms.

1. Introduction

Soar is an architecture that is to be capable of general intelligence, displaying abilities in areas such as problem-solving, planning, diagnosis, learning, etc. (Laird, Newell & Rosenbloom, ??). The architecture contains a single learning mechanism, called chunking, which saves the results of processing for application to future similar situations (Laird, Rosenbloom & Newell, 1984, Laird, Rosenbloom & Newell, 1986a). Chunking was first incorporated into Soar in early 1984 as a particular learning mechanism we had been exploring in the context of human practice.¹ Almost immediately thereafter chunking proved itself capable of several different types of learning (Laird, Rosenbloom & Newell, 1984, Rosenbloom, Laird, McDermott, Newell & Orciuch, 1985). At first tentatively, but then with increasing commitment, we adopted the hypothesis that chunking was a sufficient mechanism for all learning.

We have not engaged in a deliberate attempt to test this hypothesis in any definitive way. There currently exists no taxonomy of learning situations or processes that is well-founded enough to support such an endeavor, although partial classifications abound, such as learning by experience, learning by example, and learning by instruction. Instead we have proceeded to explore the operation of Soar in many different directions, observing the way learning entered into the various tasks and mechanisms studied.

The variety of learning exhibited by Soar seems extensive enough by now to make a listing of all the various

¹More details may be found in (Laird, Rosenbloom & Newell, 1986b).

examples worthwhile. That is the objective of this paper. First, we very briefly describe the structure of Soar. Then we list the different types of learning, using the outline of Figure 1-1. By necessity this listing treats each of the examples of in its own terms. Furthermore, it isolates them from the larger project context in which they have occurred (and indeed many parts of the total Soar project receive no mention in this list). We then redescribe this pool of instances from a more systematic standpoint. Even here the categories we use are essentially ad hoc, reflecting more our attempt to understand the diversity of what chunking has done in Soar than any a priori theoretical view of learning.

- Effects of chunking in different tasks**
 - Puzzles and toy problems
 - Syllogisms
 - Balance beam (children's development)
 - Seibel's task (stimulus-response)
 - R1-Soar (VAX computer configuration)
 - Cypress-Soar (algorithm design)
- Task-independent applications of chunking**
 - Task acquisition
 - Data chunking
 - Macro-operators
 - Constraint compilation
 - Learning from outside guidance
 - Abstraction planning
 - Explanation-based generalization
- Theoretical analysis of chunking**
 - Sources of overgeneralization in Soar
 - Meta-levels in Soar

Figure 1-1: Summary of Soar research relevant to learning

2. The Soar architecture

The following is a brief description of how Soar solves problems and learns during the process. More complete descriptions of Soar are available in (Laird, 1986) and (Laird, Newell & Rosenbloom, ??).

Soar is based on the hypothesis that all goal-directed cognitive activity can be represented as search in a problem space. A problem space is defined by a set of states, and a set of operators to move from state to state. Given a task, represented as a set of problem spaces, Soar uses knowledge to select problem spaces, states, and operators to move towards a desired state. This knowledge is represented as productions, which are syntactically similar to Ops5 (Forgy, 1981) rules. When the knowledge directly available in a given context is insufficient to determine the next thing to do immediately, Soar sets up new contexts in which it attempts to obtain the necessary knowledge.

Soar operates in terms of *decision cycles*. The decisions relate to progressing towards a desired state: selecting problem spaces to work in, states to proceed from, and operators to apply. Several productions may fire in accumulating the knowledge to make a given decision.

The situations in which Soar is unable to make progress on some goal because of inadequate directly available knowledge are called *impasses*. The types of impasses that may arise in Soar systems are determined by the architecture. For example, a common impasse, *operator-tie*, occurs when several operators are proposed as acceptable for application to a given state, and there is insufficient knowledge to choose between them. Soar exclusively sets up subgoals to resolve such impasses. In order to work on a subgoal, Soar applies all its

problem solving machinery just as it did for the higher-level goals (selecting a problem space to work on, etc.). The subgoal to resolve an operator-tie impasse is satisfied when the system acquires knowledge indicating that one of the operators initially causing the tie is now preferable to all other candidates.

When Soar finishes working on a subgoal, it can learn from its experience by building productions called *chunks* for use in future problem solving. The conditions of a chunk are the features of the pre-impasse situation that were tested during work on the subgoal to resolve the impasse. The actions of a chunk are the attributes added to the situation that are connected to the embedding supercontext; some of these attributes terminated the impasse. At first glance one would expect chunking to yield nothing more than rote learning, but generalization does occur because chunks test only relevant attributes of the problem-solving context (Laird, Rosenbloom & Newell, 1984).

3. Soar research relevant to learning

This section describes the problem, approach, and results obtained for each topic listed in Figure 1-1.

3.1. Effects of chunking in different tasks

Puzzles and toy problems: The first demonstrations of the power of chunking as a learning mechanism were obtained with tasks such as solving the eight-puzzle and playing tic-tac-toe (Laird, Rosenbloom & Newell, 1984). These results include within-trial transfer, across-trial transfer, across-task transfer, and strategy acquisition. One interesting feature is that Soar was able to learn search control knowledge that generalized across symmetric board positions, even though the representation of the tasks did not include explicit knowledge of symmetry. Since then, Soar systems have been implemented to solve many of the standard AI toy problems, and learning behavior has been studied in nearly all of these systems.

Syllogisms: Mental models have been proposed as an alternative to deduction as an account for human reasoning performance. Our work on building mental models within Soar has centered on the task of solving syllogisms: given two premises, what conclusions necessarily follow? For example, given that some B are A, and no C are B, is it necessarily true that some A are not C? Psychological studies of people performing this task show that syllogisms vary widely in difficulty. People don't necessarily make logically valid inferences; indeed, in some syllogisms, they nearly always make mistakes. Johnson-Laird (Johnson-Laird, 1983) has developed a theory that says people solve them by building mental models of possible worlds that could be described by the premises, and checking them to see if the conclusion holds. He claims the syllogisms that are the hardest for people to solve are those for which it is necessary to construct the most models. We are in the midst of an effort to map Johnson-Laird's theories into Soar. Chunking plays a strong role in this mapping, in terms of learning to draw valid conclusions. Once Soar learns a chunk stating a conclusion is valid in one model, that production will fire in any other model where that conclusion is valid.

Balance beam: Much of the learning research in Soar centers on the kind of learning that humans display in the short- to medium-term, which raises an interesting question: Can *long term* learning, i.e. the developmental behavior that children exhibit over several years, also be modeled by chunking, or are different architectural mechanisms required? In order to discover what role Soar may play in explaining the development of cognitive abilities, we are building models of problem-solving for a task studied by Siegler (Siegler, 1978). Given a beam balanced on a fulcrum at the center, the problem is to determine which side of the beam will go down when certain weights are placed at varying distances on either side. Children's abilities to solve this problem go through four fairly well-defined stages in the rules used for making predictions. Initially, they base their predictions only on the weights. Then they consider the distances in order to break ties on the weights, then

they break ties by combining the weights and distances in some fashion, and finally they use the weights and distances to compute the correct torques. We are investigating the hypothesis that chunking can be used to model this gradual increase in sophistication in solving this task. The Soar model of the balance beam task maps development onto a shift in the problem spaces used by the subject. Initially Soar has a set of attributes about the domain (such as weight and distance), but knows neither exactly which attributes are relevant, nor how to connect the attributes to the operators that predict which side of the balance beam will go down. Soar makes predictions, some of which are shown to be wrong by observation, and uses the results as feedback to create new problem spaces that make better predictions.

Seibel: As part of an investigation into the effects of practice, we are examining a reaction-time task investigated by Seibel (Seibel, 1963). The task is a psychological experiment in which the subject faces a display of ten lights, and rests his or her fingers on a set of ten buttons; during each trial, some of the lights will go on, and the subject is to press the buttons corresponding to the lights that are on as quickly as possible. The reaction time data for this task obey a power law, and earlier work showed that the power law could also be obtained from a computer simulation of the process based on chunking (Rosenbloom & Newell, 1983). Soar's chunking mechanism has a mode in which it only builds chunks for terminal subgoals, which we call *bottom-up chunking*, and which can account for the effects of practice under certain assumptions. In recent work, we have replicated the power law effect for the Seibel task in Soar, under the assumption that reaction times correspond to number of decision cycles to solution.

R1-Soar: R1-Soar was the first demonstration that Soar can combine general problem-solving methods with large amounts of domain knowledge to solve complex problems (Rosenbloom, Laird, McDermott, Newell & Orciuch, 1985). The system had about 25% of the functionality of R1, an expert system for configuring Vaxes. Because the chunking mechanism caches subgoal results during knowledge-intensive as well as knowledge-lean problem-solving, we were able to perform several experiments on the effects of learning on performance for the R1-Soar task. The knowledge gained by chunking was used in R1-Soar for search control and operator implementation. Chunking decreases the amount of effort required in problem-solving even when the system already has some search control knowledge (a factor of 3.6 for the version without search control, a factor of 2.2 for partial search control, and a factor of 1.7 for full search control). We refer to this reduction as *within-trial transfer*, because the system acquires knowledge early on that it may apply to solve later subproblems within the same trial. As would be expected from an experience-learner, knowledge transferred across trials was also substantial, yielding a solution length reduction of between a factor of 20 and 200 in subsequent runs of R1-Soar on the same task. There turned out to be little effect from transfer across different tasks.

Cypress-Soar: Another knowledge-intensive task to which we have applied Soar is algorithm design (Steier, 1987). Again, we took the approach of reimplementing a previously built system; in this case, Doug Smith's semi-automatic algorithm designer, Cypress (Smith, 1985). We set about replicating the performance of Cypress in the design of three sorting algorithms: insertion-sort, quicksort, and mergesort. Cypress makes design choices for these algorithms by instantiating a functional programming template for divide-and-conquer algorithms, and propagating constraints resulting from these instantiations by a special method of deduction known as antecedent derivation. We focused only on the methods for making the design choices in Cypress-Soar, and chose not to reimplement the antecedent derivation engine. As in R1-Soar, Cypress-Soar was able to successfully perform its tasks with varying amounts of search control knowledge. Using chunking, Cypress-Soar acquires both search control and operator implementation knowledge. The effect of within-trial transfer in one run with minimal search control was to reduce the number of decision cycles by almost 70%. Across-trial transfer on design of the same algorithm gave reductions of 90-97%. The operator implementation chunks transferred across designs of different algorithms, reducing the number of decision cycles for subsequent designs by as much as 20%. We expect the 20% across-task transfer figure would have been higher had we

implemented a Soar-based deduction engine.

3.2. Task-independent applications of chunking

Task acquisition: Currently, the standard way for Soar to acquire new tasks is for a "Soarware engineer" to analyze the task and write a set of productions implementing the problem spaces necessary for performing the task. Since Soar is intended to function autonomously, we need to define mechanisms for Soar to acquire new task spaces on its own. Our intended approach used a problem space for building problem spaces: one that would compile task descriptions directly into productions through deliberate processing (i.e., not through chunking). We have now moved to an approach that uses of the combined learning and problem solving abilities of Soar. Task acquisition is implemented in a two-phase process. The first phase is *communication* in which Soar comprehends the input and store it internally. The second phase is *interpretation* in which Soar interpretively executes the task description, and builds chunks to directly implement the task. In the communication phase, the task descriptions are input in pseudo-natural language in terms of operators, desired, initial, and illegal states. These descriptions are understood by a series of word-at-a-time comprehension operators, in a method based on Steven Small's Word Expert Parser (Small80, 1980). These operators either place information about expectations and meanings of words into the state or package a set of input tokens into higher level units. The next phase evokes a problem space for interpretation in the event of impasses resulting from attempts to execute the task. After resolving the interpretation impasses, Soar will build task-implementation chunks, so that the interpreter will not have to be evoked on subsequent task trials. These chunks turn out to be nearly identical to the ones that a programmer would have written by hand to directly implement the task. This two-phase process has now enabled Soar to acquire the eight-puzzle and missionaries-and-cannibals task spaces.

Data chunking: The evidence accumulated to date clearly demonstrates that chunking can speed up Soar's performance. However, until recently there has been no evidence that chunking could also be used to acquire new factual knowledge, and doubts have been raised as to whether it is possible for such an experience-based learning mechanism to do so. Our work on data chunking is an attempt to demonstrate how chunking can be used to acquire new knowledge. In the initial phase we are focusing on two simple memory tasks: learning to recognize, and learning to recall new objects (Rosenbloom, Laird & Newell, 1987). Recognition involves determining whether the system has ever seen the object before. Recall involves the ability to generate a representation of the new object on demand. For recognition, a chunk must be learned which contains a representation of the new object in its conditions, while for recall, the object must be represented in the actions. For both tasks, the general approach we have taken is to understand what type of problem solving is required in order to achieve the appropriate chunks. For recall, the more interesting of the two, performance is based on a generation problem space in which any potential object can be generated. This is a constructive space in which new objects are built up out of known pieces. When the system receives an object that it should learn to recall, it generates a representation of it, thus learning a chunk which can generate it in the future. At recall time, Soar's reflective ability is used to distinguish between the objects that it has learned to generate and those that it could potentially generate; that is, when all of the chunks have fired, an impasse occurs which the system uses as a signal that recall should be terminated. We are currently working on two more advanced memory tasks: cued recall and paired-associate recall. In cued recall, the recall of an item is conditioned on a set of discriminating cues that can be used as the basis for building a discrimination network for retrieval. This permits retrieval of an object by partial specifications and is being used as a building block in an implementation of an EPAM-like system (Feigenbaum & Simon, 1984) for paired-associate recall.

Macro-operators: Macro-operators have been used in a variety of systems since STRIPS (Fikes, Hart & Nilsson, 1972) to improve the performance of problem-solving systems. A macro-operator is a sequence of

operators that can be treated as a single operator. Korf demonstrated that it is possible to define a table of macro-operators showing how to transform all legal states into desired states if the problem is *serially decomposable*, that is, the subgoals of the problem can be ordered so that they only depend on preceding subgoals (Korf, 1983). It is both possible and useful to build macro tables for tasks such as the eight-puzzle and Rubik's cube. Means-ends-analysis and hill-climbing don't work for these tasks, because the problem solver has to temporarily undo previously solved subgoals in order to make progress. We have replicated Korf's results in Soar by using two problem spaces: one for the domain operations of the "conventional" task decomposition, and one for operators corresponding to serially-decomposable goals (Laird, Rosenbloom & Newell, 1986a, Laird, Newell & Rosenbloom, ??). The sets of search control chunks learned in the process of implementing each operator in the second space correspond to macro-operators. In a situation identical to Korf's macro-operator learner, in which there is no transfer between macro-operators, it would take 230 productions to encode the entire macro table (35 macro-operators) this way. However, Soar needs only 61 productions to encode the table because of the implicit generalization performed by chunking. The fine-grained encoding of the macro-operators as sets of chunks and the choice of an appropriate task representation makes this possible. More specifically, the generality arises because the operators are indifferent to the tiles not yet in place, and the chunks are invariant with respect to reflection, rotation, translation, and position within the solution path. In fact, so much transfer occurs with this representation that Soar can learn the entire macro table for the eight-puzzle in only three trials (given the correct set of problems). Similar results were obtained for the Towers of Hanoi puzzle. Soar was able to encode the entire macro table (six macro-operators) in eight chunks after a single trial.

Constraint compilation: Many tasks that have been studied in artificial intelligence can be viewed as constraint satisfaction problems (Steele, 1980). Examples are cryptarithmic, eight-queens, and the Waltz line labeling task (Waltz, 1975). If one views the satisfaction of a constraint as the application of an operator, then formulating constraint satisfaction as problem space search in Soar is not difficult. When a consistent assignment of values to all variables has been found, then the problem has been solved. This formulation leads to an interesting application for chunking: networks of constraints can be satisfied by subgoalizing in such a way that chunks can be learned for the efficient satisfaction of a constraint network, or *macro-constraints*. We have demonstrated this form of constraint compilation for small networks of digital logic elements, with the chunks being exactly the productions that would be built by hand to compile the macro-constraint.

Learning from outside guidance: We have built a system that learns search control knowledge using outside guidance in the domain of solving simple algebraic equations in one unknown (Golding, Rosenbloom & Laird, 1987). The system has a set of algebraic operators, such as *add* and *commute*, and must decide which of these operators to apply and how to instantiate them at each step of the solution. When the system has insufficient knowledge to select an operator from a set of candidates, it asks for help, which can be given to it in one of two forms: either it can accept direct advice, or it can accept a simpler problem to solve that will illustrate the principle it is to use in solving the harder problem. If direct advice is given, the system will first verify that the advice is correct, thus protecting itself from erroneous advice. During the verification, Soar builds chunks containing the relevant search control knowledge for future use. This particular type of learning also occurs in the learning apprentice systems, such as LEAP (Mitchell, Mahadevan & Steinberg, 1985), where it is called *verification-based learning* (VBL) (Mahadevan85, 1985). If the advice takes the form of a simpler problem to solve, Soar attempts to do so, either by brute-force search, or by accepting direct advice about the simpler problem. In the process, Soar will build chunks that summarize the solution, which will transfer directly to the harder problem, provided that problem selected was appropriate. This is a simple form of analogy; we have not yet investigated more complex analogies that would attempt to adapt the simpler problem solution to the harder problem with deliberate processing.

Abstraction planning: Another kind of problem-solving behavior that has traditionally been obtained in AI systems (such as ABSTRIPS (Sacerdoti, 1974)) through special-purpose mechanisms is abstraction planning. Chunking in Soar has already been used to achieve the effects of non-abstraction planning: ties in operator selection impasses are often resolved by lookahead, and the chunks formed during the lookahead apply to reduce effort in problem-solving based on the result of the lookahead. One way to make planning even more cost-effective is to use abstracted problem spaces for the lookahead search (Unruh, Rosenbloom & Laird, 1987). Abstraction of problem spaces in Soar is not based on the explicit manipulation of a declarative specification of operators. Rather, the abstraction is a consequence of ignoring parts of the problem space as needed during problem-solving. The particular abstractions possible are determined by the manner in which the problem space has been decomposed into independent subcomponents such as initial state creation, operator precondition testing, operator implementation, state evaluation and goal testing. When searching in an abstracted problem space, only the productions implementing the operations that are part of the abstracted space fire, and the chunks summarizing the processing in the abstracted search transfer to the non-abstracted space. The abstracted problem-solving has paid attention to a subset of what it would normally examine. Therefore, the chunks learned can be more general than those learned normally, and the appropriateness of the abstraction determines how useful the chunks are. This was tested in R1-Soar with several abstractions. The abstractions used were removing the backplane-box operator, removing the modules-backplane operator, and removing the state information about module width. Two abstractions together saved 75% of decision cycles, learning alone saved 60%, and learning together with the two abstractions saved 77%. Issues currently being investigated in this work are the automatic creation and selection of appropriate abstractions, and the use of multiple levels of abstractions.

Explanation-based generalization: Using the framework for EBG described by Mitchell, Keller and Kedar-Cabelli (Mitchell, Keller, Kedar-Cabelli, 1986), we have shown that chunking in Soar implements a variant of explanation-based generalization (Rosenbloom & Laird, 1986). The domain theory is implemented in some Soar problem space, and an explanation is the trace of the problem-solving in this space. Chunking uses this trace in the same way as goal regression uses a proof or explanation. The aim of regression is to describe the goal concept in terms of operational predicates, where chunking produces a description of the concept in terms of predicates on attributes of the context that existed prior to the impasse. The condition finding algorithm in Soar's chunking is essentially the same as regression in EBG except it uses instantiated goals and rules rather than parametrized versions. The mapping works not only for operator implementation, but also for learning search control knowledge, where the goal concept is the preference for a selection of an operator. No extra mechanisms are needed to produce the results of EBG with chunking; the problem solving that is analogous to EBG's explanation is exactly the same as is used by Soar in other problem solving. To test this mapping, we took several examples in (Mitchell, Keller, Kedar-Cabelli, 1986) and implemented problem spaces in Soar that acquired nearly identical concepts to the ones acquired by "conventional" EBG algorithms.

3.3. Theoretical analysis of chunking

Sources of overgeneralization in Soar: Soar exhibits overgeneralization while chunking in several instances. We now believe that we understand most, if not all, sources of overgeneralization in Soar (Laird, Rosenbloom & Newell, 1986c). Our investigation was based on the theory of chunking as *knowledge compilation*, that is, the conversion of knowledge from one form to another so that it may be used in similar situations but more efficiently than before. For this conversion to be correct and avoid overgeneralization in Soar, all of the relevant aspects of the pre-impasse situation that lead to the creation of results must be included as conditions of a chunk. This is easy as long as the results of a subgoal depend only on the productions that fired during the subgoal. However, in some subgoals, the appropriate result (such as success or failure) is most easily detected by testing some feature of the problem solving process, such as the available operators are

exhausted and there nothing more to do. For example, if the highest number of a set is to be found, and each step of the problem solving involves comparing one of the numbers to the highest found so far, the problem is finished when no other numbers are left to be compared. In Soar, detecting that there is nothing more to do means testing that no more productions can fire for the current situation which in Soar becomes testing for the existence of an impasse. For chunking to be correct, it must determine why no more productions fired and includes these reasons in the chunk. Although this may be logically possible, it is computationally intractable and the resulting chunks can be less efficient than the original problem solving. Therefore, these conditions are not included and the resulting chunks are overgeneral. This problem arises whenever the problem solving for a goal is dependent on its own process state (meta-information) which is made available in Soar though the impasses. Because of these problems, we are currently investigating a variety of methods to recover from overgeneralization in these cases when it cannot be prevented.

Chunking and meta-levels: We have also analyzed Soar in terms of concepts such as meta-levels, introspection and reflection (Rosenbloom, Laird & Newell, 1986). Several meta-levels were identified, and chunking was cast as a means for shifting knowledge from the problem-space meta-level, where its application is slow and deliberate, to the production meta-level, where its application is fast and automatic. This production meta-level acts as a long-term cache for the results of problem solving at the higher problem space level, and improves the efficiency of the problem solver by reducing the need to climb the hierarchy of meta-levels as often.

4. Soar as a learning architecture

Figure 4-1 lists some questions that might be asked about a particular learning architecture in order to assess its abilities and compare it to others. This section attempts to answer these questions with regard to Soar.

- When can the architecture learn?
 - On what kinds of tasks?
 - On what occasions during performance of the tasks?
- From what can the architecture learn?
 - Internal sources of knowledge?
 - External sources of knowledge?
- What can the architecture learn?
 - What kinds of procedural knowledge?
 - What kinds of declarative knowledge?
- When can the architecture apply learned knowledge?
 - Within the same trial
 - Across different trials
 - Across different tasks

Figure 4-1: Questions to ask about learning architectures

4.1. When can Soar learn?

When working on different types of tasks: These tasks span the range from simple and knowledge-lean to complex and knowledge-intensive, from routine to design and discovery. Soar has exhibited interesting learning on instances of all of these, and is theoretically capable of learning on any task it can perform. The simple tasks include the eight-puzzle, tic-tac-toe, towers of Hanoi, missionaries-and-cannibals, solving simple

algebraic equations in one unknown, satisfying logical constraint networks, solving logical syllogisms, balance beam, simple blocks world, eight-queens, monkey-and-bananas, and the water jug problems. The complex tasks include computer configuration, algorithm design, and applying sociological theories. Routine tasks include paired-associate learning and the Seibel task. One discovery task is algorithm design.

Whenever a result is generated: The chunking mechanism is the same for results at any level in the goal hierarchy, and is invoked independently of the success or failure of the goal. For example, R1-Soar builds a chunk for the top-level goal during the successful configuration of a VAX backplane. Learning from failure occurs during the missionaries-and-cannibals runs; if a move is made that results in there being more cannibals than missionaries on a bank or in the boat (an illegal state), Soar learns not to make that move again.

4.2. From what can Soar learn?

Internal sources of knowledge: The normal source of learning for Soar is its own problem solving experience, since it usually functions autonomously. For example, Soar has learned to solve the missionaries-and-cannibals problems by brute-force search, and to satisfy macro-constraints based on solutions to constraint networks found previously. The mechanism for all this is uniform. The chunking mechanism takes the traces of the problem solving as input. The productions that led to the creation of all results are backtraced. The conditions of those productions testing elements that existed prior to the impasse are variabilized and made part of the actions of the chunk. Any elements created during the subgoal that are tied to the superstate are considered part of the result, and corresponding actions are made part of the right-hand side of the chunk.

External sources of knowledge: Several of the research efforts described earlier illustrate how Soar can learn from knowledge provided to it externally. In the algebraic domain, the user can supply direct advice on which of several operators to apply to an equation. Soar can also obtain the same knowledge if the user supplies it with a similar, simpler problem to solve. The chunks built after solving the simpler problem will transfer to provide the search control for the original problem. Two other sources of external knowledge now being tapped by Soar are pseudo-natural language descriptions of task problem spaces and declarative facts.

4.3. What can Soar learn?

Procedural knowledge: A step in a procedure in Soar corresponds to the replacement of a single element in a goal context. A goal context consists of a goal, problem space, state and operator. We can analyze what procedural knowledge Soar can learn, and has been demonstrated to learn, by examining what aspects of this context replacement are subject to modification by chunking.

- **Goals:** The set of possible goal types forms a complete set, fixed by the architecture; goal types are determined by the decision cycle in response to impasses in a manner also fixed by the architecture. Consequently, the effect of chunking on the goal element of a context is limited to avoiding the need for subgoal.
- **Problem spaces:** A problem space consists of a set of possible states and operators, and the architecture does not impose any restrictions on what constitutes a legal set. Therefore, Soar must generate candidate problem spaces, and select from among the candidates. For generation, Soar can either acquire problem spaces from external descriptions, or can modify problem spaces it already has to obtain new ones. The work on task acquisition has shown Soar to be capable of the former. For the latter, problem spaces have been abstracted from old ones by omitting operators in the work on abstraction planning, and Soar builds new problem spaces during the balance beam task as well. There is nothing preventing Soar from learning problem space selection knowledge, we have only a few examples of learning this knowledge, in the context of task acquisition and the balance beam task.
- **States:** The selection of a problem space imposes certain constraints on the creation and selection

of states. Our main experience to date with acquiring initial state creation and selection knowledge has been through task acquisition. For creation of states other than initial ones, the processing, and thus the knowledge to acquire by chunking, is primarily determined by the operator selected for application to the previous state. Much of the transfer due to chunking comes from learning how to create states, or implement operators, without subgoal. For example, in Cypress-Soar, an operator in the top-level space creates a divide-and-conquer sorting algorithm to satisfy a sorting specification. The first time through the problem, this operator will be implemented by subgoaling into a special space for designing divide-and-conquer algorithms, and a chunk will be built that creates an algorithm such as quicksort for that specification. On subsequent trials, that chunk will fire to create the new state with the algorithm directly.

- **Operators:** Operator creation is part of task acquisition, and an equivalent effect was achieved in learning macro tables for the eight-puzzle and Tower of Hanoi. For operator selection knowledge, examples abound in the toy problems, in R1-Soar and Cypress-Soar and in the research on learning from outside guidance and abstraction planning. Most operator selection, or search control, chunks result from summarizing the processing in lookahead search, where it is learned that a particular operator application will be on the path to a solution.

Declarative knowledge: The work on recognition, free recall, and cued recall shows that Soar can learn declarative knowledge.

4.4. When can Soar apply learned knowledge?

Within the same trial: Knowledge that is learned while solving a problem applied to improve performance in a later stage of solving the same problem is called *within-trial transfer*. This may be measured by comparing the number of decision cycles Soar takes to solve a given problem without learning against how many it takes with learning. The reduction can be substantial: a factor of 2 for some runs of the eight puzzle, over a factor of 2 for Cypress-Soar, and a factor of 3 for R1-Soar. The transfer occurs because of the presence of repeated subgoals during a trial for these problems.

Across different trials: When the system builds chunks for every terminated goal, the only effort required for solving a problem on the second and subsequent trials is the processing demand by solving the problem with no search. A chunk in R1-Soar reduces the number of decision cycles in a solution of the same problem on subsequent trials by a factor of 20 to 200 in some cases. If the system is learning bottom-up, where chunks are only built for terminal goals, the processing that is repeated most often gets chunked first. Eventually, the results of bottom-up chunking will be the same as those of all-goals chunking. Work on the Seibel reaction-time task has shown that the across-trial transfer with bottom-up chunking leads to a power law effect on performance.

Across different tasks: Across-task transfer occurs in Soar when similar subgoals show up in different tasks; this happened in Cypress-Soar, in the Seibel task, in several of the toy problems, and to some extent in R1-Soar. For example, in the divide-and-conquer formulation of all the sorting algorithms synthesized by Cypress-Soar, a sub-algorithm must be created to sort lists of length zero or one. Once this sub-algorithm, which merely returns its input, was designed for one algorithm, the chunk built during that design could fire on the design of other algorithms. Once one of the three algorithms had been designed, such chunks reduced by 10 to 20 percent the decision cycles needed to design the other two algorithms.

4.5. What variety of learning has Soar not yet exhibited?

In order to support the claim that chunking is a truly general learning mechanism, it must be studied further in a number of areas. A partial list of these areas follows:

- **Very large problems:** We have not yet tested learning in Soar in large (greater than 10^3 rules) systems for solving very difficult problems, though this may soon be possible as a consequence of work in progress on Designer-Soar (for algorithm design tasks) and Neomycin-Soar (for medical diagnosis tasks).
- **Learning from exhaustion:** We are currently investigating issues of learning from exhaustion, that is, learning when a goal terminates because Soar runs out of things to do. This is difficult because the reason for the exhaustion is not an explicit input to the chunking mechanism.
- **Complex analogies:** In learning from outside guidance provided as a simpler problem to solve, Soar performs a rudimentary form of analogy, which requires that the chunks transfer directly. Soar is not yet capable of more complex analogical reasoning involving deliberate processing to adapt the solution to the simpler problem. Related to this is the problem of how to use multiple examples.
- **Representation shifts:** In the category of problem space creation, we have yet to demonstrate the creation of "completely new" representations in a task. While there is a moderate shift of representation in the balance beam task, it is still the case that the set of possible problem space modifications must be somehow "designed in" beforehand.
- **Across-space transfer:** Soar has not yet exhibited transfer of knowledge across different problem spaces. Demonstration of such transfer would significantly increase the transfer of knowledge gained during the performance of one task to other tasks.

5. Conclusion

We have described how Soar successfully learns in a wide range of tasks and situations. Most of these results could previously only be obtained from special-purpose learning mechanisms. The variety of learning possible in Soar does not arise from collecting all these mechanisms into one system. Rather, they have been produced with a single mechanism, which is general enough to apply to all problem-solving behavior. Thus the diversity of problem-solving of which Soar is capable is directly responsible for the varieties of learning in Soar.

Acknowledgements

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under contracts F33615-81-K-1539 and N00039-83-C-0136. Additional partial support was provided by the Sloan Foundation and some computing support was supplied by the SUMEX-AIM facility (NIH grant number RR-00785). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Sloan Foundation, the National Institute of Health, or the U.S. Government.

References

- Feigenbaum, E. A. and Simon, H. A. EPAM-like models of recognition and learning. *Cognitive Science*, 1984, 8, 305-336.
- Fikes, R., Hart, P., and Nilsson, N. Learning and executing generalized robot plans. *Artificial Intelligence*, 1972, 3, 251-288.
- Forgy, C. L. *OPSS User's Manual* (Tech. Rep. CMU-CS-81-135). Carnegie-Mellon University, Computer Science Department, July 1981.
- Golding, A., Rosenbloom, P. S., and Laird, J. E. Learning general search control from outside guidance. In *Proceedings of the Tenth International Conference on Artificial Intelligence*. Milano, Italy: Proceedings of IJCAI-87, 1987. To appear.
- Johnson-Laird, P. N. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press, 1983.
- Korf, R. E. *Learning to solve problems by searching for macro-operators*. Doctoral dissertation, Carnegie-Mellon University, 1983.
- Laird, J. E. *Soar user's manual* (Tech. Rep. ISL-15). Xerox Corporation, January 1986.
- Laird, J. E., Newell, A., and Rosenbloom, P. S. Soar: An architecture for general intelligence. (*Artificial Intelligence*, in press).
- Laird, J. E., Rosenbloom, P. S., and Newell, A. Towards chunking as a general learning mechanism. In *Proceedings of the Fourth National Conference on Artificial Intelligence*. Austin, Texas: The American Association for Artificial Intelligence, 1984.
- Laird, J. E., Rosenbloom, P. S., and Newell, A. Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, 1986, Vol. 1(1).
- Laird, J. E., Rosenbloom, P. S., and Newell, A. *Universal subgoalting and chunking: The automatic generation and learning of goal hierarchies*. Hingham, MA: Kluwer, 1986.
- Laird, J. E., Rosenbloom, P. S., and Newell, A. Overgeneralization during knowledge compilation in Soar. In *Proceedings of the Workshop on Knowledge Compilation*. Otter Crest, OR: Oregon State University, 1986.
- Mahadevan, S. Verification-based learning: A generalization strategy for inferring problem reduction methods. In *Proceedings of the Ninth International Conference on Artificial Intelligence*. Los Angeles, CA: Proceedings of IJCAI-85, 1985.
- Mitchell, T. M., Keller, R. M., and Kedar-Cabelli, S. T. . Explanation-based generalization: A unifying view. *Machine Learning*, 1986, 1(1), .
- Mitchell, T. M., Mahadevan, S. and Steinberg, L. I. . LEAP: A learning apprentice for VLSI design. In *Proceedings of the Ninth International Conference on Artificial Intelligence*. Los Angeles, CA: Proceedings of IJCAI-85, 1985.
- Rosenbloom, P. S., and Laird, J. E. Mapping explanation-based generalization onto Soar. In *Proceedings of the Fifth National Conference on Artificial Intelligence*. Philadelphia, PA: American Association for Artificial Intelligence, 1986.
- Rosenbloom, P.S. and Newell, A. The chunking of goal hierarchies: A generalized model of practice. In R. S. Michalski (Ed.), *Proceedings of the International Machine Learning Workshop*. Department of Computer Science, University of Illinois, Urbana-Champaign, IL, 1983.
- Rosenbloom, P. S., Laird, J. E., and Newell, A. Meta-levels in Soar. In *Preprints of the Workshop on Meta-level Architectures and Reflection*. Sardinia: , 1986.
- Rosenbloom, P. S., Laird, J. E. and Newell, A. Knowledge-level learning in Soar. In *Proceedings of AAAI-87*. American Association for Artificial Intelligence, 1987. To appear.

- Rosenbloom, P. S., Laird, J. E., McDermott, J., Newell, A., and Orciuch, E. R1-Soar: An experiment in knowledge-intensive programming in a problem-solving architecture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1985, 7(5), 561-569.
- Sacerdoti, E. D. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 1974, 5, 115-135.
- Seibel, R. Discrimination reaction time for a 1,023-alternative task. *Journal of Experimental Psychology*, 1963, 66(3), 215-226.
- Siegler, R. S. The origin of scientific reasoning. In R. S. Siegler (Ed.), *Children's thinking: What develops?*], undef [R. S. Siegler (Ed.), Siegler, R. S. *The origin of scientific reasoning*.])Hillsdale, N.J.: Lawrence Erlbaum and Associates, 1978.
- Small, S. *Word expert parsing: a distributed theory of word-based natural language understanding* (Tech. Rep. TR-954). Department of Computer Science, University of Maryland, 1980. (PhD thesis).
- Smith, D. R. Top-down synthesis of divide-and-conquer algorithms. *Artificial Intelligence*, 1985, 27(1), 43-96.
- Steele, G. L. Jr. *The Definition and Implementation of a Computer Programming Language Based on Constraints* (Tech. Rep. AI-TR 595). MIT, August 1980.
- Steier, D. M. Cypress-Soar: A case study in search and learning in algorithm design. In *Proceedings of the Tenth International Conference on Artificial Intelligence*. Milano, Italy: Proceedings of UCAI-87, 1987. To appear.
- Unruh, A., Rosenbloom, P. S., and Laird, J. E. Dynamic abstraction problem solving in Soar. 1987. In preparation.
- . Understanding line drawings of scenes with shadows. In P. Winston (Ed.), *The Psychology of Computer Vision*], undef [P. Winston (Ed.). *Understanding line drawings of scenes with shadows*.])New York: McGraw-Hill, 1975.